

XP-002180985

P.D. 31-07-1999	18
P. 1-18	

## Data-Over-Cable Service Interface Specifications

### Radio Frequency Interface Specification

(abstracts)

**SP-RF1v1.1-102-990731**

**INTERIM  
SPECIFICATION**

#### Notice

This document is a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry in general. Neither CableLabs nor any member company is responsible for any liability of any nature whatsoever resulting from or arising out of use or reliance upon this specification by any party. This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding its accuracy, completeness, or fitness for a particular purpose. Distribution of this document is restricted pursuant to the terms of separate access agreements negotiated with each of the parties to whom this document has been furnished.

© Copyright 1999 Cable Television Laboratories, Inc.  
All rights reserved.

**BEST AVAILABLE COPY**

## Document Status Sheet

**Document Control Number:** SP-RFiv1.1-I02-990731

**Reference:** Radio Frequency Interface Specification

**Revision History:** I01 — First Interim Release, March 11, 1999  
I02 — Second Interim Release, July 31, 1999

**Date:** July 31, 1999

**Status Code:** ~~Work in Process~~      ~~Draft~~      ~~Interim~~      ~~Released~~

**Distribution Restrictions:** ~~Cable Labs only~~      ~~CL Reviewers~~      ~~CL Vendor~~      ~~Public~~

## Key to Document Status Codes

- Work in Process** An incomplete document, designed to guide discussion and generate feedback, that may include several alternative requirements for consideration.
- Draft** A document in specification format considered largely complete, but lacking review by cable industry and vendors. Drafts are susceptible to substantial change during the review process.
- Interim** A document which has undergone rigorous cable industry and vendor review, suitable for use by vendors to design in conformance with, and suitable for field testing.
- Released** A stable document, reviewed, tested and validated, suitable to enable cross-vendor interoperability.

## Contents

<b>1</b>	<b>SCOPE AND PURPOSE.....</b>	<b>1</b>
1.1	SCOPE.....	1
1.2	REQUIREMENTS.....	1
1.3	BACKGROUND.....	1
1.3.1	<i>Service Goals.....</i>	<i>1</i>
1.3.2	<i>Reference Architecture.....</i>	<i>2</i>
1.3.3	<i>Categories of Interface Specification.....</i>	<i>4</i>
1.3.4	<i>Statement of Compatibility.....</i>	<i>5</i>
<b>2</b>	<b>FUNCTIONAL ASSUMPTIONS.....</b>	<b>7</b>
2.1	BROADBAND ACCESS NETWORK.....	7
2.2	EQUIPMENT ASSUMPTIONS.....	7
2.2.1	<i>Frequency Plan.....</i>	<i>7</i>
2.2.2	<i>Compatibility with Other Services.....</i>	<i>7</i>
2.2.3	<i>Fault Isolation Impact on Other Users.....</i>	<i>8</i>
2.2.4	<i>Cable System Terminal Devices.....</i>	<i>8</i>
2.3	RF CHANNEL ASSUMPTIONS.....	8
2.3.1	<i>Transmission Downstream.....</i>	<i>8</i>
2.3.2	<i>Transmission Upstream.....</i>	<i>9</i>
2.4	TRANSMISSION LEVELS.....	10
2.5	FREQUENCY INVERSION.....	10
<b>3</b>	<b>COMMUNICATION PROTOCOLS.....</b>	<b>11</b>
3.1	PROTOCOL STACK.....	11
3.1.1	<i>CM and CMTS as Hosts.....</i>	<i>11</i>
3.1.2	<i>Data Forwarding Through the CM and CMTS.....</i>	<i>12</i>
3.2	THE MAC FORWARDER.....	15
3.2.1	<i>Rules for Data-Link-Layer Forwarding.....</i>	<i>16</i>
3.3	NETWORK LAYER.....	16
3.3.1	<i>Requirements for IGMP Management.....</i>	<i>17</i>
3.4	ABOVE THE NETWORK LAYER.....	18
3.5	DATA LINK LAYER.....	18
3.5.1	<i>LLC Sublayer.....</i>	<i>19</i>
3.5.2	<i>Link-Layer Security Sublayer.....</i>	<i>19</i>
3.5.3	<i>MAC Sublayer.....</i>	<i>19</i>
3.6	PHYSICAL LAYER.....	19
3.6.1	<i>Downstream Transmission Convergence Sublayer.....</i>	<i>19</i>
3.6.2	<i>PMD Sublayer.....</i>	<i>20</i>
<b>4</b>	<b>PHYSICAL MEDIA DEPENDENT SUBLAYER SPECIFICATION.....</b>	<b>21</b>
4.1	SCOPE.....	21
4.2	UPSTREAM.....	21

4.2.1	Overview.....	21
4.2.2	Modulation Formats.....	22
4.2.3	FEC Encode.....	25
4.2.4	Scrambler (Randomizer) .....	25
4.2.5	Preamble Prepend.....	26
4.2.6	Burst Profiles.....	26
4.2.7	Burst Timing Convention.....	29
4.2.8	Transmit Power Requirements .....	30
4.2.9	Fidelity Requirements.....	31
4.2.10	Frame Structure.....	35
4.2.11	Signal Processing Requirements .....	36
4.2.12	Upstream Demodulator Input Power Characteristics .....	37
4.2.13	Upstream Electrical Output from the CM.....	38
4.3	DOWNSTREAM.....	38
4.3.1	Downstream Protocol.....	38
4.3.2	Scalable Interleaving to Support Low Latency .....	38
4.3.3	Downstream Frequency Plan.....	39
4.3.4	CMTS Output Electrical.....	39
4.3.5	Downstream Electrical Input to CM.....	40
4.3.6	CM BER Performance.....	40
4.3.7	CMTS Timestamp Jitter.....	41
5	DOWNSTREAM TRANSMISSION CONVERGENCE SUBLAYER .....	43
5.1	INTRODUCTION.....	43
5.2	MPEG PACKET FORMAT .....	43
5.3	MPEG HEADER FOR DOCSIS DATA-OVER-CABLE.....	44
5.4	MPEG PAYLOAD FOR DOCSIS DATA-OVER-CABLE.....	44
5.5	INTERACTION WITH THE MAC SUBLAYER .....	45
5.6	INTERACTION WITH THE PHYSICAL LAYER.....	46
5.7	MPEG HEADER SYNCHRONIZATION AND RECOVERY .....	46
6	MEDIA ACCESS CONTROL SPECIFICATION.....	47
6.1	INTRODUCTION.....	47
6.1.1	Overview.....	47
6.1.2	Definitions .....	47
6.1.3	Future Use.....	49
6.2	MAC FRAME FORMATS .....	49
6.2.1	Generic MAC Frame Format.....	49
6.2.2	Packet-Based MAC Frames .....	53
6.2.3	ATM Cell MAC Frames.....	55
6.2.4	Reserved PDU MAC Frames.....	55
6.2.5	MAC-Specific Headers.....	56
6.2.6	Extended MAC Headers .....	61

## 8.2 Upstream Service Flow Scheduling Services

The following sections define these upstream Service Flow scheduling services: Unsolicited Grant Service (UGS), Real-Time Polling Service (rtPS), Unsolicited Grant Service with Activity Detection (UGS-AD), Non-Real-Time Polling Service (nrtPS), Best Effort (BE) service, and a Committed Information Rate (CIR) service.

### 8.2.1 Unsolicited Grant Service

The intent of UGS is to reserve specific upstream transmission opportunities for specific real-time traffic flows. The CMTS MUST provide fixed size data grants at periodic intervals to the Service Flow. The CM MUST use only unsolicited data grants for upstream transmission. The key service information elements are the Unsolicited Grant Size, the Nominal Grant interval, the Tolerated Grant Jitter and the Request/Transmission Policy. (Refer to Appendix M)

The Unsolicited Grant Synchronization Header (UGSH) in the Service Flow EH Element (refer to Section 6.2.6.3.2) is used to pass status information from the CM to the CMTS regarding the state of the UGS Service Flow. The most significant bit of the UGSH is the Queue Indicator (QI) bit. The CM MUST set this flag once it detects that this Service Flow has exceeded its transmit queue depth. Once the CM detects that the Service Flow's transmit queue is back within limits, it MUST clear the QI flag. The flag allows the CMTS to provide for long term compensation for conditions such as lost maps or clock rate mismatch's by issuing additional grants.

The CMTS MUST NOT grant in excess of active Grants per Interval, excluding the case when the QI bit in the UGSH is set. In this case, the CMTS MAY grant up to 1% additional bandwidth for clock rate mismatch compensation. The CMTS policing of the Service Flow remains unchanged.

### 8.2.2 Real-Time Polling Service

The intent of rtPS is to reserve upstream transmission opportunities for real-time traffic flows (such as voice over IP). These Service Flows receive periodic transmission opportunities regardless of network congestion, but these Service Flows release their transmission reservations to other Service Flows when inactive. The CMTS polls rtPS SIDs on a periodic basis (typically on the order of tens of milliseconds or less) for current upstream traffic usage through the use of unicast request opportunities.

The CMTS MUST provide periodic unicast request opportunities. The CM MUST use only unicast request opportunities in order to obtain upstream transmission opportunities (the CM MUST use unsolicited data grants for upstream transmission as well). The key service information elements are the Nominal Polling Interval, the Tolerated Poll Jitter and the Request/Transmission Policy.

### 8.2.3 Unsolicited Grant Service with Activity Detection

The intent of UGS-AD is to reserve upstream transmission opportunities for real-time traffic flows (such as voice-over-IP with silence suppression) in a way that takes advantage of the reduced latency of Unsolicited Grants, when active, but has the efficiency of Real-Time Polling when inactive.

The CMTS MUST provide periodic unicast grants, when the flow is active, but MUST revert to providing periodic unicast request opportunities when the flow is inactive. [The CMTS can detect flow inactivity by detecting unused grants. However, the algorithm for detecting a flow changing from an active to an inactive state is dependent on the CMTS implementation] The CM MUST use only unicast request opportunities in order to obtain upstream transmission opportunities. The CM MUST use unsolicited data grants for upstream transmission as well. The key service parameters are the Nominal Polling Interval, the Tolerated Poll Jitter, the Nominal Grant Interval, the Tolerated Grant Jitter, the Unsolicited Grant Size and the Request/Transmission Policy.

In UGS-AD service, when restarting UGS after an interval of RTP, the CMTS SHOULD provide additional grants in the first (and/or second) grant interval such that the CM receives a total of one grant for each grant interval from the time the CM requested restart of UGS, plus one additional grant. (Refer to Appendix M) Because the Service Flow is provisioned as a UGS flow with a specific grant interval and grant size, when restarting UGS, the CM MUST NOT request a different sized grant than the already provisioned UGS flow. As with any Service Flow, changes can only be requested with a DSC command.<sup>1</sup>

The Service Flow Extended Header Element allows for the CM to dynamically state how many grants per interval are required to support the number of flows with activity present. In UGS/AD, the CM MAY use the Queue Indicator Bit in the UGSH. The remaining seven bits of the UGSH define the Active Grants field. This field defines the number of grants within a Nominal Grant Interval that this Service Flow currently requires. The CM MAY indicate from 0 to Grants Per Interval grants active per Nominal Grant Interval. This field allows the CM to signal to the CMTS to dynamically adjust the number of grants per interval that this UGS Service Flow is actually using. The CM MUST NOT request more than the number of Grants per Interval in the ActiveQoSParameterSet.

#### 8.2.4 Non-Real-Time Polling Service

The intent of nrtPS is to set aside upstream transmission opportunities for non-real-time traffic flows (such as high bandwidth FTP). These flows receive some transmission opportunities during network congestion. The CMTS typically polls nrtPS SIDs on an (periodic or non-periodic) interval on the order of one second or less.

The CMTS MUST provide timely unicast request opportunities. The CM MUST use both contention request and unicast request opportunities in order to obtain upstream transmission opportunities (the CM MUST use unsolicited data grants for upstream transmission as well). The key service elements are Nominal Polling Interval, Reserved Minimum Traffic Rate, Maximum Sustained Traffic Rate, Request/Transmission Policy and Traffic Priority.<sup>2</sup>

#### 8.2.5 Best Effort Service

The intent of the Best Effort (BE) service is to provide efficient service to best effort traffic. With BE service, the CM MUST use all (contention and unicast) request opportunities<sup>3</sup>, and all unicast data transmission opportunities. The CM MAY use contention data opportunities as appropriate. The key service elements are the Minimum Reserved Traffic Rate, the Maximum Sustained Traffic Rate, and the Traffic Priority.<sup>4</sup>

#### 8.2.6 Other Services

##### 8.2.6.1 Committed Information Rate (CIR)

A Committed Information Rate (CIR) service can be defined a number of different ways. For example, it could be configured by using a Best Effort service with a Reserved Minimum Traffic Rate or a nrtPS with a Reserved Minimum Traffic Rate.

1. Sentenc added 06/22/99 per rfi-n-99043 ew
2. Edited 06/22/99 per rfi-n-99043 ew
3. With appropriate deference for contention request opportunities. Refer to 7.4.1.
4. Edited 06/22/99 per rfi-n-99043 ew

## 8.3 Fragmentation

Fragmentation is an upstream CM “modem capability”. The CMTS **MUST** enable or disable this capability on a per-modem basis with a TLV in the Registration Response. The per-modem basis provides compatibility with DOCSIS 1.0 CMs. Once fragmentation is enabled for a DOCSIS 1.1 modem, fragmentation is enabled on a per-Service Flow basis via the Request/Transmission Policy Configuration Settings. When enabled for a Service Flow, fragmentation is initiated by the CMTS when it grants bandwidth to a particular CM with a grant size that is smaller than the corresponding bandwidth request from the CM. This is known as a **Partial Grant**.

### 8.3.1 CM Fragmentation Support

Fragmentation is essentially encapsulation of a portion of a MAC Frame within a fixed size fragmentation header and a fragment CRC. Concatenated PDUs, as well as single PDUs, are encapsulated in the same manner. Baseline Privacy, if enabled, is performed on each fragment as opposed to the complete original MAC frame.

The CM **MUST** perform fragmentation according to the flow diagram in Figure 8-8. The phrase “untransmitted portion of packet” in the flow diagram refers to the entire MAC frame when fragmentation has not been initiated and to the remaining untransmitted portion of the original MAC frame when fragmentation has been initiated.

#### 8.3.1.1 Fragmentation Rules:

1. Any time fragmentation is enabled and the grant size is smaller than the request, the CM **MUST** fill the partial grant it receives with the maximum amount of data (fragment payload) possible accounting for fragmentation overhead and physical layer overhead.
2. The CM **MUST** send up a piggyback request any time there is no later grant or grant pending for that SID in MAPs that have been received at the CM.
3. If the CM is fragmenting a frame<sup>1</sup>, any piggyback request **MUST** be made in the BPI EHDR portion of the fragment header.
4. In calculating bandwidth requests for the remainder of the frame (concatenated frame, if concatenated) that has been fragmented, the CM **MUST** request enough bandwidth to transmit the entire remainder of the frame plus the 16-byte fragment overhead and all associated physical layer overhead.
5. If the CM does not receive a grant or grant pending within the ACK time of sending a request, the CM **MUST** backoff and re-request for the untransmitted portion of the frame until the bandwidth is granted or the CM exceeds its retry threshold.
6. If the CM exceeds its retry threshold while requesting bandwidth, the CM discards whatever portion of the frame was not previously transmitted.
7. The CM **MUST** set the F bit and clear the L bit in the first fragment of a frame.
8. The CM **MUST** clear the F and L bits in the fragment header for any fragments that occur between the first and last fragments of a frame.
9. The CM **MUST** set the L bit and clear the F bit in the last fragment of a frame.
10. The CM **MUST** increment the fragment sequence number sequentially for each fragment of a frame transmitted.
11. If a frame is to be encrypted and the frame is fragmented, the frame is encrypted only at the fragment layer with encryption beginning immediately after the fragment header HCS and continuing through the fragment CRC.
12. Frames sent in immediate data (request/data) regions **MUST NOT** be fragmented.

---

1. Note, ‘frame’ always refers to either frames with a single Packet PDU or concatenated frames.

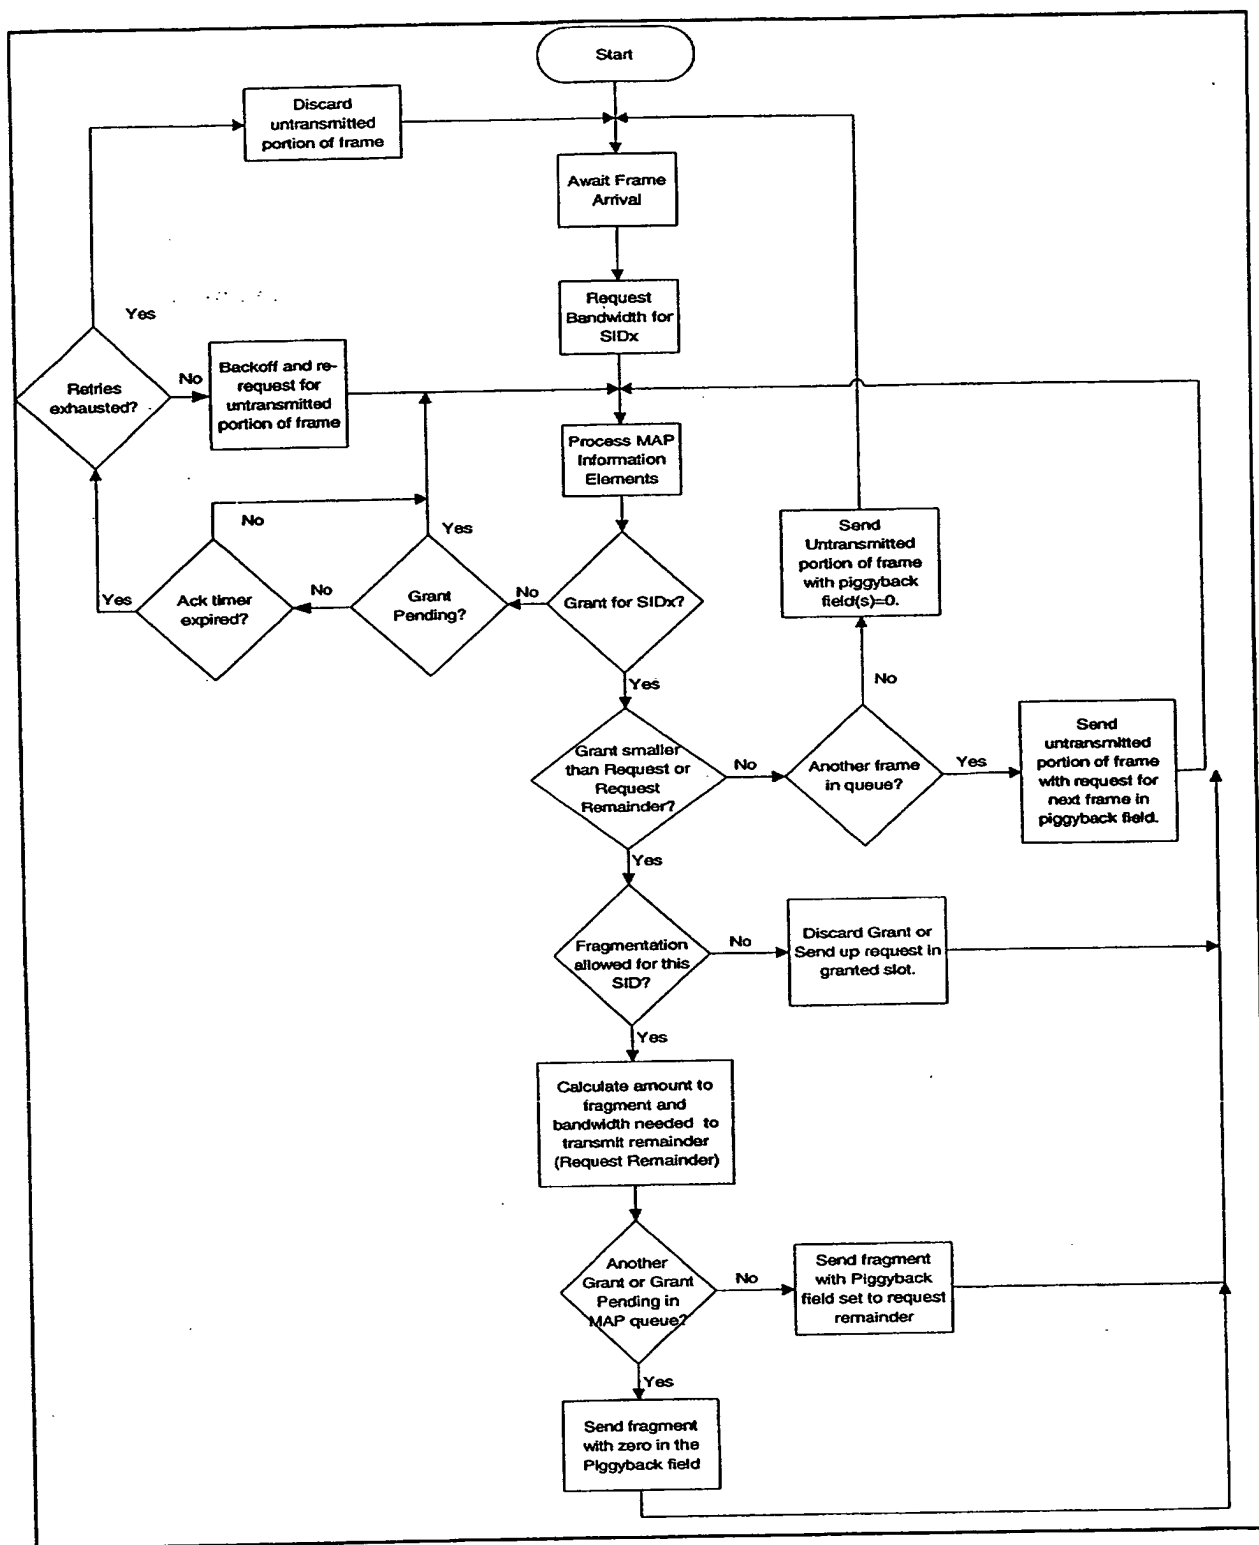


Figure 8-8. CM Fragmentation Flowchart



### 8.3.2 CMTS Fragmentation Support

At the CMTS, the fragment is processed similarly to an ordinary packet with the exception that the baseline privacy encryption starts just after the fragmentation header as opposed to being offset by 12 bytes.

The CMTS has two modes it can use to perform fragmentation. The Multiple Grant Mode assumes that the CMTS retains the state of the fragmentation. This mode allows the CMTS to have multiple partial grants outstanding for any given SID. The Piggybacking Mode assumes the CMTS does NOT retain any fragmentation state. Only one partial grant is outstanding, so that the CM inserts the remaining amount into the Piggyback field of the fragment header. The type of mode being used is determined by the CMTS. In all cases, the CM operates with a consistent set of rules.

#### 8.3.2.1 Multiple Grant Mode

A CMTS MAY support Multiple Grant Mode for performing fragmentation.

Multiple Grant Mode allows the CMTS to break a request up into two or more grants in a single or over successive maps and it calculates the additional overhead required in the remaining partial grants to satisfy the request. In Multiple Grant Mode, if the CMTS cannot grant the remainder in the current MAP, it MUST send a grant pending (zero length grant) in the current MAP and all subsequent MAPs to the CM until it can grant additional bandwidth. If there is no grant or grant pending in subsequent maps, the CM MUST re-request for the remainder. This re-request mechanism is the same as that used when a normal REQ does not receive a grant or grant pending within the ACK time.

If a CM receives a grant pending IE along with a fragment grant, it MUST NOT piggyback a request in the extended header of the fragment transmitted in that grant.

In the case where the CM misses a grant and re-requests the remaining bandwidth, the CMTS MUST recover without dropping the frame.

Due to the imprecision of the mini-slot to byte conversion process, the CMTS MUST make sure that any fragment payload remainder is greater than a mini-slot (i.e. the imprecision amount). Failure to do this may cause the CMTS to issue a grant that is not needed as the CM has completed transmission of the fragment payload remainder using a previous partial grant.<sup>1</sup> This may cause the CM to get out of sync with the CMTS by inadvertently starting a new fragmentation.

#### 8.3.2.2 Piggyback Mode

A CMTS MAY support Piggyback Mode for performing fragmentation.

If the CMTS does not put another partial grant or a grant pending in the MAP in which it initiates fragmentation on a SID, the CM MUST automatically piggyback for the remainder. The CM calculates how much of a frame can be sent in the granted bandwidth and forms a fragment to send it. The CM utilizes the piggyback field in the fragment extended header to request the bandwidth necessary to transfer the remainder of the frame. Since the CMTS did not indicate a multiple grant in the first fragment MAP, the CM MUST keep track of the remainder to send. The request length, including physical-layer and fragmentation overhead, for the remainder of the original frame is inserted into the piggyback request byte in the fragmentation header.

If the fragment HCS is correct, the piggybacked request, if present, is passed on to the bandwidth allocation process while the fragment itself is enqueued for reassembly. Once the complete MAC Frame is reassembled,

---

1. Sentence clarified 06/22/99 per rfi-n-99043 ew

any non-privacy extended headers are processed if the packet HCS is correct, and the packet is forwarded to the appropriate destination.

### 8.3.3 Fragmentation Example

#### 8.3.3.1 Single Packet Fragmentation

Refer to Figure 8-8. Assume that fragmentation has been enabled for a given SID.

1. (Requesting State)- CM wants to transmit a 1018 byte packet. CM calculates how much physical layer overhead (POH) is required and requests the appropriate number of minislots. CM makes a request in a contention region. Go to step 2.
2. (Waiting for Grant)- CM monitors MAPs for a grant or grant pending for this SID. If the CM's ACK time expires before the CM receives a grant or grant pending, the CM retries requesting for the packet until the retry count is exhausted - then the CM gives up on that packet. Go to step 3.
3. (First Fragment)- Prior to giving up in step 2, the CM sees a grant for this SID that is less than the requested number of minislots. The CM calculates how much MAC information can be sent in the granted number of minislots using the specified burst profile. In the example in Figure 8-9, the first grant can hold 900 bytes after subtracting the POH. Since the fragment overhead (FRAG HDR, FHCS, and FCRC) is 16 bytes, 884 bytes of the original packet can be carried in the fragment. The CM creates a fragment composed of the FRAG HDR, FHCS, 884 bytes of the original packet, and an FCRC. The CM marks the fragment as first and prepares to send the fragment. Go to step 4.
4. (First Fragment, multiple grant mode)- CM looks to see if there are any other grants or grant pendings enqueued for this SID. If so, the CM sends the fragment with the piggyback field in the FRAG HDR set to zero and awaits the time of the subsequent grant to roll around. --- go to step 6. If there are not any grants or grant pendings, go to step 5.
5. (First Fragment, piggyback mode)- If there are no other grants or grant pendings for this SID in this MAP, the CM calculates how many minislots are required to send the remainder of the fragmented packet, including the fragmentation overhead, and physical layer overhead, and inserts this amount into the piggyback field of the FRAG HDR. The CM then sends the fragment and starts its ACK timer for the piggyback request. In the example in Figure 8-9, the CM sends up a request for enough minislots to hold the POH plus 150 bytes (1018-884+16). Go to step 6.
6. (Waiting for Grant)- The CM is now waiting for a grant for the next fragment. If the CM's ACK timer expires while waiting on this grant, the CM should send up a request for enough minislots to send the remainder of the fragmented packet, including the fragmentation overhead, and physical layer overhead. Go to step 7.
7. (Receives next fragment grant)- Prior to giving up in step 6, the CM sees another grant for this SID. The CM checks to see if the grant size is large enough to hold the remainder of the fragmented packet, including the fragmentation overhead and physical layer overhead. If so, go to step 10. If not, go to step 8.
8. (Middle Fragment, multiple grant mode)- Since the remainder of the packet (plus overhead) will not fit in the grant, the CM calculates what portion will fit. The CM encapsulates this portion of the packet as a middle fragment. The CM then looks for any other grants or grant pendings enqueued for this SID. If either are present, the CM sends the fragment with the piggyback field in the FRAG HDR set to zero and awaits the time of the subsequent grant to roll around. --- go to step 6. If there are not any grants or grant pendings, go to step 9.
9. (Middle Fragment, piggyback mode) - The CM calculates how many minislots are required to send the remainder of the fragmented packet, including the fragmentation overhead and physical layer overhead, and inserts this amount into the piggyback field of the FRAG HDR. The CM then sends the fragment and starts its ACK timer for the piggyback request. Go to step 6.

10. (Last Fragment) - The CM encapsulates the remainder of the packet as a last fragment. If there is no other packet enqueued or there is another grant or a grant pending enqueued for this SID, the CM places a zero in the REQ field of the FRAG HDR. If there is another packet enqueued with no grant or grant pending, the CM calculates the number of minislots required to send the next packet and places this number in the REQ field in the FRAG HDR. The CM then transmits the packet. Go to step 11. In the example in Figure 8-9, the grant is large enough to hold the remaining 150 bytes plus POH.
11. (Normal operation)- The CM then returns the normal operation of waiting for grants and requesting for packets. If at any time fragmentation is enabled and a grant arrives that is smaller than the request, the fragmentation process starts again as in step 2.

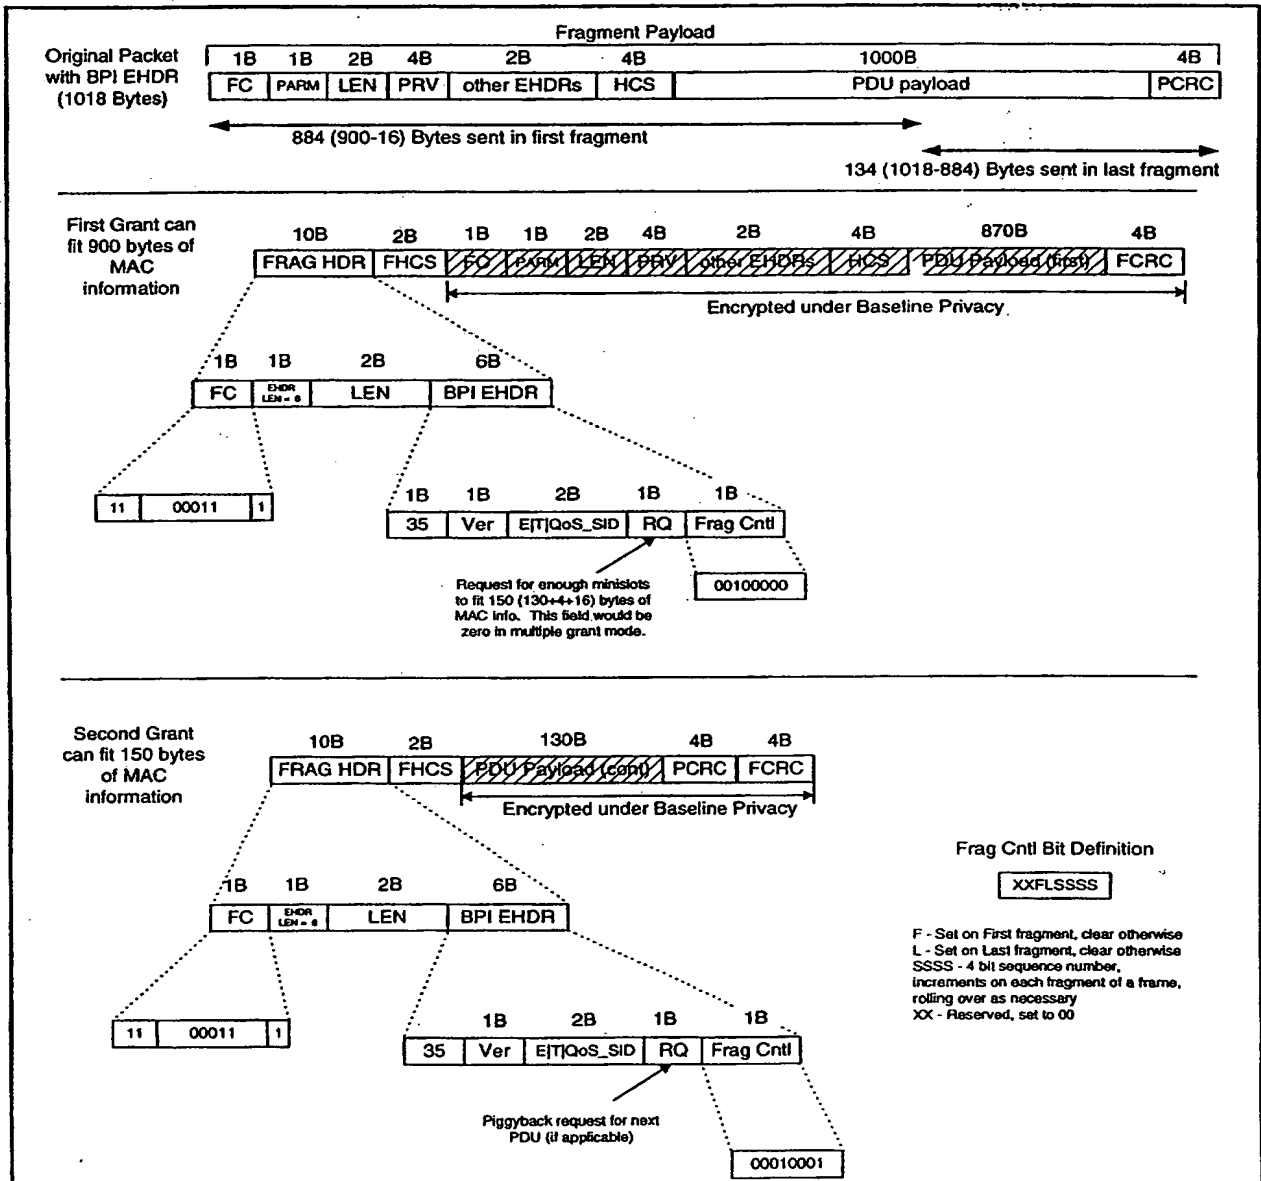


Figure 8-9. Example of Fragmenting a Single Packet

### 8.3.3.2 Concatenated Packet Fragmentation

After the CM creates the concatenated packet, the CM treats the concatenated packet as a single PDU. Figure 8-10 shows an example of a concatenated packet broken into 3 fragments. Note that the packet is fragmented without regard to the packet boundaries within the concatenated packet.

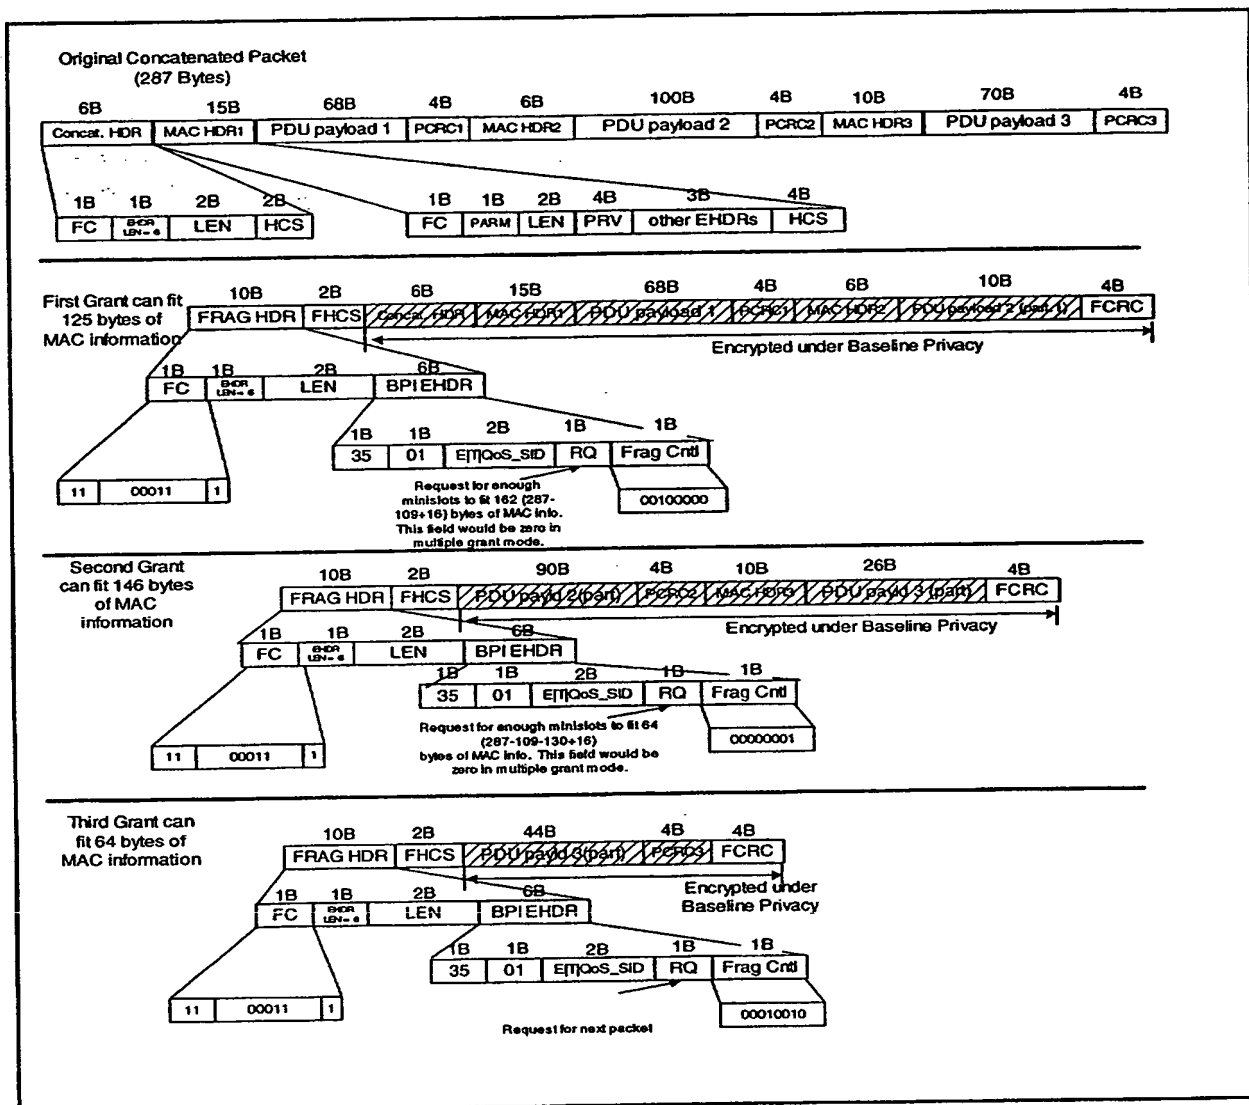


Figure 8-10. Fragmented Concatenated Packet Example

## Appendix M. Unsolicited Grant Services<sup>1</sup>

This appendix discusses the intended use of the Unsolicited Grant Service (UGS) and Unsolicited Grant Service with Activity Detection (UGS-AD) and includes specific examples.

### M.1 Unsolicited Grant Service (UGS)

#### M.1.1 Introduction

Unsolicited Grant Service is an Upstream Flow Scheduling Service Type that is used for mapping constant bit rate (CBR) traffic onto Service Flows. Since the upstream is scheduled bandwidth, a CBR service can be established by the CMTS scheduling a steady stream of grants. These are referred to as unsolicited because the bandwidth is predetermined, and there are no ongoing requests being made.

The classic example of a CBR application of interest is Voice over Internet Protocol (VoIP) packets. Other applications are likely to exist as well.

Upstream Flow Scheduling Services are associated with Service Flows, each of which is associated with a single Service ID (SID). Each Service Flow may have multiple Classifiers. Each Classifier may be associated with a unique CBR media stream. Classifiers may be added and removed from a Service Flow. Thus, the semantics of UGS must accommodate single or multiple CBR media streams per SID.

For the discussion within this Appendix, a Subflow will be defined as the output of a Classifier. Since a VoIP session is identified with a Classifier, a Subflow in this context refers to a VoIP session.

#### M.1.2 Configuration Parameters

- Nominal Grant Interval
- Unsolicited Grant Size
- Tolerated Grant Jitter
- Grants per Interval

Explanation of these parameters and their default values are provided in Appendix C.

#### M.1.3 Operation

When a Service Flow is provisioned for UGS, the Nominal Grant Interval is chosen to equal the packet interval of the CBR application. For example, VoIP applications with 10 ms packet sizes will require a Nominal Grant Interval of 10 ms. The size of the grant is chosen to satisfy the bandwidth requirements of the CBR application and relates directly to the length of the packet.

When multiple Subflows are assigned to a UGS service, multiple grants per interval are issued. There is no explicit mapping of Subflows to grants. The multiple grants per interval form a pool of grants in which any subflow can use any grant.

It is assumed in this operational example the default UGS case of no concatenation and no fragmentation.

---

1. Section renamed per rfi-n-99043, item b. ew

### M.1.4 Jitter

Figure M-1 shows the relationship between Grant Interval and Tolerated Grant Jitter, and shows an example of jitter on subflows.

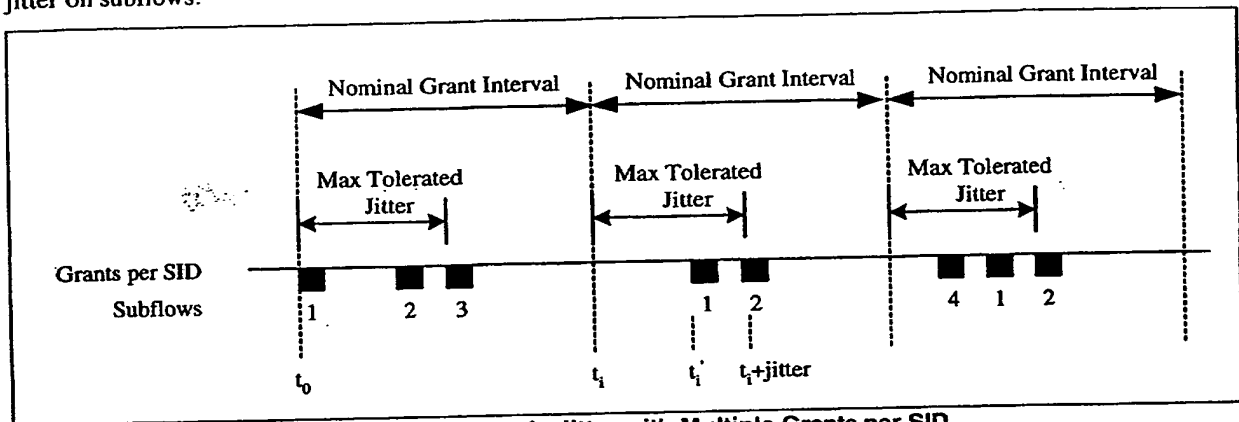


Figure M-1. Example Jitter with Multiple Grants per SID

For only one Grant per Interval, the Tolerated Grant Jitter is the maximum difference between the actual grant time ( $t_i'$ ) and the nominal grant time ( $t_i$ ). For multiple Grants per Interval, the Tolerated Grant Jitter is the maximum difference between the actual time of the last grant in the group of grants and the nominal grant time ( $t_i$ ). If the arrival of any grant is at  $t_i'$ , then  $t_i \leq t_i' \leq t_i + \text{jitter}$ .

Figure M-1 demonstrates how a Subflow will be jittered even though the individual grants may not move from their relative position. During the first interval, three VoIP sessions are established, and they happen fall on the three grants. In the second interval, VoIP session 3 has been torn down. Since the CMTS does not know which Subflow is associated with which grant, it decides to remove the first grant. The remaining two calls shift to the other two grants. In the third interval, a new VoIP session 4 and a new grant have been added. The new call happens to fall on the new grant. The net effect is that the Subflows may move around within their jitter interval.

The advantage of a small jitter interval is that the VoIP receive jitter buffer may be kept small. The disadvantage is that this places a scheduling constraint on the CMTS.

The boundary of a Nominal Grant Interval is arbitrary and is not communicated between the CMTS and the CM.

**Note:** More dramatic events like the loss of a downstream MAP, or the frequency hopping of an upstream may cause subflows to jitter outside of this jitter window.

### M.1.5 Synchronization Issues<sup>1</sup>

There are two synchronization problems that occur when carrying CBR traffic such as VoIP sessions across a network. The first is a frequency mismatch between the source clock and the destination clock. This is managed by the VoIP application, and is beyond the scope of this specification. The second is the frequency mismatch between the CBR source/sinks, and the bearer channel that carries them.

Specifically, if the clock that generates the VoIP packets towards the upstream is not synchronized with the clock at the CMTS which is providing the UGS service, the VoIP packets may begin to accumulate in the CM. This could also occur if a MAP was lost, causing packets to accumulate.

1. Title changed per rfi-r-99043 06/21/99 ew

When the CM detects this condition, it asserts the Queue Indicator in the Service Flow EH Element. The CMTS will respond by issuing an occasional extra grant so as to not exceed 1% of the provisioned bandwidth. (This corresponds to a maximum of one extra grant every one hundred grants). The CMTS will continue to supply this extra bandwidth until the CM deasserts this bit.

A similar problem occurs in the downstream. The far end transmitting source may not be frequency synchronized to the clock which drives the CMTS. Thus the CMTS should police at a rate slightly higher than the exact provisioned rate to allow for this mismatch and to prevent delay buildup or packet drops at the CMTS.

## **M.2 Unsolicited Grant Service with Activity Detection (UGS-AD)**

### **M.2.1 Introduction**

Unsolicited Grant Service with Activity Detection (UGS-AD) is an Upstream Flow Scheduling Service Type. This section describes one application of UGS-AD which is the support for Voice Activity Detection (VAD). VAD is also known as Silence Suppression and is a voice technique in which the transmitting CODEC sends voice samples only when there is significant voice energy present. The receiving CODEC will compensate for the silence intervals by inserting silence or comfort noise equal to the perceived background noise of the conversation.

The advantage of VAD is the reduction of network bandwidth required for a conversation. It is estimated that 60% of a voice conversation is silence. With that silence removed, that would allow a network to handle substantially more traffic.

Subflows in this context will be described as active and inactive. Both of these states of within the MAC Layer QOS state known as Active.

### **M.2.2 MAC Configuration Parameters**

The configuration parameters include all of the normal UGS parameters, plus:

- Nominal Polling Interval
- Tolerated Poll Jitter

Explanation of these parameters and their default values are provided in <Appendix C>.

### **M.2.3 Operation**

When there is no activity, the CMTS sends polled requests to the CM. When there is activity, the CMTS sends Unsolicited Grants to the CM. The CM indicates in the UGS\_Parm field of the Service Flow EH Element in each packet of each Unsolicited Grant the number of Grants per Interval that it currently requires. The CM may request up to the maximum active Grants per Interval. The CM constantly sends this state information so that no explicit acknowledgment is required from the CMTS.

It is left to the implementation of the CM to determine activity levels. Implementation options include:

- Having the MAC layer service provide an activity timer per Classifier. The MAC layer service would mark a Subflow inactive if packets stopped arriving for a certain time, and mark a Subflow active the moment a new packet arrived. The number of Grants requested would equal the number of active Subflows.
- Having a higher layer service entity such as an embedded media client which indicates activity to the MAC layer service.

When the CM is receiving polled requests and it detects activity, the CM requests enough bandwidth for one Grant per Interval. If activity is for more than one Subflow, the CM will indicate this in the UGS\_Parm field of the Service Flow EH Element beginning with the first packet is sends.

When the CM is receiving Unsolicited Grants, then detects new activity, and asks for one more grant, there will be a delay in time before it receives the new grant. During that delay, packets may build up at the CM. When the new Unsolicited Grant is added, the CMTS will burst extra Grants to clear out the packet buildup.

When the CM is receiving Unsolicited Grants, then detects inactivity on a Subflow and asks for one less grant, there will be a delay in time before the reduction in Grants occurs. If there has been any build up of packets in the upstream transmit queue, the extra grants will reduce or empty the queue. This is fine, and keeps system latency low. The relationship of which Subflow is getting which specific grant will also change. This effect appears as low frequency jitter that the far end must manage.

When the CM is receiving Unsolicited Grants and detects no activity on any of its Subflows, it will send one packet with the Service Flow EH Element with the UGS\_Parm field set to zero grants, and then cease transmission. The CMTS will switch from UGS mode to Real Time Polling mode.

It is not necessary for the CMTS to separately monitor packet activity since the CM does this already. Worst case, if the CMTS misses the last packet which indicated zero grants, the CMTS and CM would be back in sync at the beginning of the next talk spurt. Because of this scenario, when the CM goes from inactive to active, the CM must be able to restart transmission with either Polled Requests or Unsolicited Grants.

#### M.2.4 Example




Figure M-2. VAD Start-Up and Stop

Figure M-2 shows an example of a single G.711 (64 kbps) voice call with a packet size of 10 ms, and a receive jitter buffer that requires a minimum of 20 ms of voice (thus 2 packets) before it will begin playback.



Assume voice begins at time zero. After a nominal processing delay and a 10 ms packetization delay, the DSP CODEC generates voice packets which are then transferred to the upstream transmit queue. The next Polled Request is used which results in the start of the Unsolicited Grants some time later. Additional Unsolicited Grants are immediately issued to clear out the upstream queue.

These packets traverse the network and arrive at the receive jitter buffer. The 20 ms minimum jitter buffer is met when the second packet arrives. Because the packets arrived close together, only an additional few milliseconds of latency has been added. After a nominal processing delay, playout begins.

When the voice spurt ends, the CM sends one remaining packet with no payload, and with the Service Flow EH Element with the UGS\_Parm field set to zero grants. Some time later, UGS stops, and Real Time Polling begins.

### M.2.5 Talk Spurt Grant Burst

The extra burst of Unsolicited Grants when a flow becomes active is necessary because the jitter buffer at the receiving CODEC typically waits to have a minimum amount of voice samples before beginning the playout. Any delay between the arrival of these initial packets will add to the final latency of the phone call. Thus, the sooner the CMTS recognizes that the CM has packets to send and can empty the CM's buffer, the sooner those packet will reach the receiver, and the lower the latency that will be incurred in the phone call.

It is an indeterminate problem as to how many grants must be burst. When the CM makes its request for an additional grant, one voice packet has already accumulated. The CM has no idea how many extra grants to request as it has no idea of the round trip response time it will receive from the CMTS, and thus how many packets may accumulate. The CMTS has a better idea, although it does not know the far end jitter buffer requirements.

The solution is for the CMTS to choose the burst size, and burst these grants close together at the beginning of the talk spurt. This occurs when moving from Real Time Polling to UGS, and when increasing the number of UGS Grants per Interval.

A typical start-up latency that will be introduced by the Request to Grant response time is shown in Table M-1.

Variable		Example Value	
1.	The time taken from when the voice packet was created to the time that voice packet arrives in the CM upstream queue.	0 - 1	ms
2.	The time until a polled request is received. The worst case time is the Polled Request Interval.	0 - 5	ms
3.	The Request-Grant response time of the CMTS. This value is affected by MAP length and the number of outstanding MAPS.	5 - 15	ms
4.	The round trip delay of the HFC plant including the downstream interleaving delay.	1 - 5	ms
<b>Total</b>		<b>6 - 26</b>	<b>ms</b>

Table M-1. Example Request to Grant Response Time

This number will vary between CMTS implementations, but a reasonable number of extra grants to expect from the example above would be:

UGS Interval	Extra Grants for New Talk Spurts
10 ms	2
20 ms	1
30 ms	0

Table M-2. Example Extra Grants for New Talk Spurts

Once again it is worth noting that the CMTS and CM cannot and do not associate individual Subflows with individual grants. That means that when current Subflows are active and a new Subflow becomes active, the new Subflow will immediately begin to use the existing pool of grants. This potentially reduces the start up latency of new talk spurts, but increases the latency of the other Subflows. When the burst of grants arrives, it is shared with all the Subflows, and restores or even reduces the original latency. This is a jitter component. The more Subflows that are active, the less impact that adding a new Subflow has.

### M.2.6 Admission Considerations

Note that when configuring the CMTS admission control, the following factors must be taken into account.

VAD allows the upstream to be over provisioned. For example, an upstream that might normally handle 24 VoIP sessions might be over provisioned as high as 36 (50%) or even 48 (100%). Whenever there is over provisioning, there exists the statistical possibility that all upstream VoIP sessions may become active. At that time, the CMTS may be unable to schedule all the VoIP traffic. Additionally, the talk spurt grant bursts would be stretched out. CM implementations of VAD should recognize this possibility, and set a limit as to how many packets they will allow to accumulate on its queue.

Occasional saturation of the upstream during VAD can be eliminated by provisioning the maximum number of permitted VoIP sessions to be less than the maximum capacity of the upstream with all voice traffic (24 in the previous example). VAD would cause the channel usage to drop from 100% to around 40% for voice, allowing the remaining 60% to be used for data and maintenance traffic.

**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

**BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☒ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☒ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** \_\_\_\_\_

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.**

**THIS PAGE BLANK (USPTO)**